

Spatial Soft-Core Caching

Derya Malak and Muriel Médard

Research Laboratory of Electronics, MIT, Cambridge, MA, USA
Email: {deryam, medard}@mit.edu

Edmund M. Yeh

Northeastern University, Boston, MA, USA
Email: eyeh@ece.neu.edu

Abstract—We propose a decentralized spatial soft-core cache placement (SSCC) policy for wireless networks. SSCC yields a spatially balanced sampling via negative dependence across caches, and can be tuned to satisfy cache size constraints with high probability. Given a desired cache hit probability, we compare the 95% confidence intervals of the required cache sizes for independent placement, hard-core placement and SSCC policies. We demonstrate that in terms of the required cache storage size, SSCC can provide up to more than 180% and 100% gains with respect to the independent and hard-core placement policies, respectively. SSCC can be used to enable proximity-based applications such as device-to-device communications and peer-to-peer networking as it promotes the item diversity and reciprocation among the nodes.

I. INTRODUCTION

Distributed caching is a powerful technique to minimize the total average delay by replacing the backhaul capacity with storage capacity at small cells [1], and to enable spectral reuse and throughput gain in networks [2]. The goal of an efficient cache placement is to maximize the hit probability, i.e. the probability of obtaining the desired item from a neighboring cache. This is affected by the demand distribution, network topology, range of communication, and cache storage size.

Fundamental limits of caching have been studied in [2], in which the content placement phase is carefully designed so that a single coded multicast transmission can satisfy different demands. Capacity scaling laws have been explored in [2], and rate-memory and storage-latency tradeoffs have been studied in [3]. Caching has been studied in the context of device-to-device (D2D) communications in [4], and interference management in [5], and in optimization of cloud and edge processing for radio access networks in [6].

Temporal caching models have been analyzed in [7] for popular cache replacement algorithms, e.g. least recently used (LRU), least-frequently used, and most recently used cache update. Decentralized spatial LRU caching strategies have been developed in [8]. These combine the temporal and spatial aspects of caching, and approach the performance of centralized policies as the coverage increases. However, they are restricted to the LRU principle. A time-to-live (TTL) policy with a stochastic capacity constraint and low variance has been proposed in [9]. The BitTorrent protocol employs the rarest first and choke algorithms to promote diversity of the pieces among peers, and foster reciprocation, respectively. These have been demonstrated in the context of peer-to-peer (P2P) file replication in the Internet [10]. A good piece replication algorithm should minimize the time spent in the transient state.

There exist studies focusing on decentralized (geographic) content placement policies such as [1], [11], [12], [13]. The main focus of the literature in this direction is to maximize the average cache hit probability subject to an average cache constraint. This optimization problem can be solved as a convex program. However, to the best of our knowledge, the related literature does not provide guarantees in terms of (a) how far-off the average cache size is from reality, (b) how far-off the average cache hit rate is from reality, and (c) how stable the cache hit probability across the caches.

In the current paper, we provide a decentralized spatial soft-core cache placement (SSCC) policy. Since the cache storage size is finite, it is intuitive to have an exclusion range-based caching model such that the caches storing the same item are never closer to each other than some given distance (negative dependence), so as to promote diversity and reciprocation. SSCC roots in spatially balanced sampling, which is motivated by the request arrivals. For example, in P2P networking, the actual demand distribution is not known by nodes, and the cache updates in each peer are triggered by the requests. Furthermore, the traffic density in cellular networks is in general not uniform across the network, and the peak hour density can be approximated by a log-normal distribution [14]. Hence, instead of having a fixed exclusion range, it is desirable to have a variable exclusion range, depending on the popularity of the item. The SSCC policy come to the fore by putting a mark distribution on the exclusion range of an item based on its popularity. The marks may correspond to the detection ranges or the transmit powers of the nodes in heterogeneous network scenarios. Our objective is to address the issues (a)-(c) above in order to provide a better trade-off between the actual cache hit rate and the cache size violation probability. Our main contributions and use cases of SSCC are:

- i. SSCC has desirable properties: spatially balanced sampling across caches, concentration of the cache size, better cache over-provisioning, and multi-hop connectivity.
- ii. SSCC yields a better cache hit probability-cache violation probability tradeoff than the state of the art. In terms of the required cache storage size, SSCC can provide more than 180% and 100% gains with respect to independent placement [11], and hard-core placement [13], respectively.
- iii. SSCC is suited for enabling proximity-based applications (D2D, P2P), and offloading mobile users in networks.
- iv. SSCC has connections with rarest first caching as it promotes the item diversity and reciprocation among the nodes. Hence, it can be well-suited for P2P applications.

Notation. Let Φ denote the mother point process (p.p.), and Φ_{th} be the child p.p. obtained via the thinning of Φ . Let π be a spatial caching policy that yields a set of child p.p.'s $\{\Phi_{th,i}\}_i$, where $\Phi_{th,i}$ is the set of retained points that cache item i . Let A be a given bounded convex set in \mathbb{R}^2 containing the origin, and rA be its dilation by the factor r . $\mathbb{1}\{A\}$ is the indicator of event A . Let B be a bounded Borel set. Let $\Phi(B)$ be the random number of points of the spatial p.p. Φ which lie in B . Any receiver can obtain the desired item i if it is within a critical communication range R_c . Assume that $B = B_0(R_c)$, where $B_0(r)$ is a ball in \mathbb{R}^2 with radius r , centered at origin.

II. HOW TO OPTIMIZE THE CACHING GAIN

The locations of the nodes (caches) in the network are modeled by a homogeneous Poisson point process (PPP) Φ in \mathbb{R}^2 with intensity λ . There are M items in the network, each having the same size, and each node has the same cache storage size $N < M$. Each user makes requests based on a Zipf popularity distribution over the set of the items. The probability mass function (pmf) of such requests (demand) is given by $p_r(i) = i^{-\gamma_r} / \sum_{j=1}^M j^{-\gamma_r}$, where γ_r determines the tilt of the Zipf distribution. The demand profile is the Independent Reference Model (IRM), i.e., the standard synthetic traffic model in which the request distribution does not change over time [15]. The request distribution is uniform across the network, i.e., isotropic, and does not change over time. Hence, the intensity of the requests for item i , i.e. λ_i , is proportional to its demand probability $p_r(i)$. Let $\mathcal{I} \sim p_r$ be the random variable that models the demand. Each node is associated with the variables $z_{xi} = \mathbb{1}\{i \in \text{Cache}(x)\}$ that denote whether item i is available in its cache or not. There is also a cost w_k associated with obtaining an item within the presence of k nodes in the range. Given these parameters, consider the caching gain function of the following form:

$$F(Z) = \mathbb{E}_{\mathcal{I}} \left[\sum_{k=1}^{\infty} w_k \left(1 - \prod_{k'=1}^k (1 - z_{p_{k'}\mathcal{I}}) \right) \right], \quad (1)$$

where (1) can be used to model multi-hop coverage scenarios as in [12], and Boolean Model coverage scenarios as in [11], [13]. Let $\lambda_i = p_r(i)$, and $w_k = \mathbb{P}(\Phi(B) = k)$, which is the probability that k caches (nodes of the original p.p. Φ) cover the typical receiver, and w_0 is the probability of having no connection. Assume that k^* is the first index such that a transmitter has the desired item i . Then, from (1), the caching gain for item i is $\sum_{k=k^*}^{\infty} w_k = \mathbb{P}(\Phi(B) \geq k^*)$, which is the same as probability of having at least k^* transmitters. Equivalently, the cost of caching is $\sum_{k=1}^{k^*-1} w_k$.

Since both the multi-hop and Boolean coverage scenarios are equivalent up to scaling, we focus on the second scenario.

We have the following immediate observation.

Proposition 1. $F(Z)$ is convex if $z_{p_{k'}\mathcal{I}}$'s are negatively associated (NA) [16] across $k' \in \{1, \dots, k\}$, for all $i \in \{1, \dots, M\}$.

Proof. Exploiting (1), we have the following relation:

$$\mathbb{E}[F(Z)] = \mathbb{E}_{\mathcal{I}} \left[\sum_{k=0}^{\infty} w_k \left(1 - \mathbb{E} \left[\prod_{k'=1}^k (1 - z_{p_{k'}\mathcal{I}}) \right] \right) \right] \stackrel{(a)}{\geq}$$

$$\mathbb{E}_{\mathcal{I}} \left[\sum_{k=0}^{\infty} w_k \left(1 - \prod_{k'=1}^k (1 - \mathbb{E}[z_{p_{k'}\mathcal{I}}]) \right) \right] = F(\mathbb{E}[Z]),$$

where (a) is due to that $\mathbb{E} \left[\prod_{k'=1}^k (1 - z_{p_{k'}\mathcal{I}}) \right] \leq \prod_{k'=1}^k (1 - \mathbb{E}[z_{p_{k'}\mathcal{I}}])$ as $z_{p_{k'}\mathcal{I}}$'s are NA across $k' \in \{1, \dots, k\}$, $\forall i$. \square

From Prop. 1, $\mathbb{E}[F(Z)] \geq F(\mathbb{E}[Z])$. The expected cache hit probability obtained via NA placement upper bounds the independent placement solution with probabilities $\mathbb{E}[z_{p_{k'}\mathcal{I}}]$. NA has desirable properties in terms of sampling and concentration. Some important results that hold for independent variables, e.g., the Chernoff-Hoeffding bounds, and the Kolmogorov's inequality [16], also hold for NA variables.

From Prop. 1, it is clear that in terms of average cache hit performance, NA placement performs better than independent placement. Therefore, our main focus is on a class of placement policies that are NA. We also demonstrate that NA placement policies have lower variance across the nodes, hence are more stable than independent placement policies.

III. A SOFT-CORE CACHING MODEL

The spatial soft-core caching (SSCC) policy is constructed from the underlying PPP Φ by removing certain nodes depending on the positions of the neighboring nodes, and on the marks and weights attached to them. It generalizes the Matérn II hard-core p.p. (MatII) such that there is a distinct distribution modeling the exclusion radius of each item.

For each item i , let $\tilde{\Phi}_i = \{(x_k, m_k^{(i)}, v_k^{(i)})\}_k$ be a homogeneous independently marked PPP with intensity λ , and i.i.d. \mathbb{R}^2 -valued marks, where $\Phi = \{x_k\}$, and $\{(m_k^{(i)}, v_k^{(i)})\}$ is the random bivariate mark. The first component $m^{(i)}$ of the bivariate mark is referred to as mark, and has distribution $\mu^{(i)}$. The mark of item i , i.e., $m^{(i)}$, denotes its exclusion radius, and depends on its popularity in the network. If item i is more popular than item j , then $m^{(i)}$ is stochastically dominated¹ by $m^{(j)}$. The second component $v^{(i)}$ of the bivariate mark is weight, which serves as a weight in the thinning procedure, and has distribution $\nu_{m^{(i)}}^{(i)}$ which might depend on $m^{(i)}$.

Let $\Phi_{th,i}$ be a soft-core p.p. that denotes the set of points that cache item i . The cache placement model is such that item i is stored in cache $x_k \in \Phi$ if and only if cache x_k is kept as a point of $\Phi_{th,i}$. Equivalently, we have

$$z_{x_k i} = \mathbb{1}\{i \in \text{Cache}(x_k)\} = \mathbb{1}\{x_k \in \Phi_{th,i}\}. \quad (2)$$

Node x_k is retained as a point of $\Phi_{th,i}$ with probability $\mathbb{E}[z_{x_k i}] = p(x_k, m_k^{(i)}, v_k^{(i)}, \Phi)$. The weights are i.i.d. and uniformly distributed, i.e. $v_k^{(i)} \sim U[0, 1]$, for each node x_k and item i . The marks $m_k^{(i)}$ are distributed according to $\mu^{(i)}$ for each x_k , and i . For the special case of MatII, i.e., when the marks are fixed, we optimized the exclusion radii in [13].

The number of items in cache x_k is the sum of the individual items' indicator functions $C(x_k) = \sum_i \mathbb{1}\{i \in \text{Cache}(x_k)\}$. The cache size constraint has to be satisfied on average, i.e.

$$N = \mathbb{E}[C(x_k)] = \sum_i p(x_k, m_k^{(i)}, v_k^{(i)}, \Phi), \quad x_k \in \Phi. \quad (3)$$

¹ X is stochastically dominated by Y , which is denoted by $X \leq^{st} Y$, if for all increasing functions g , we have $\mathbb{E}[g(X)] \leq \mathbb{E}[g(Y)]$.

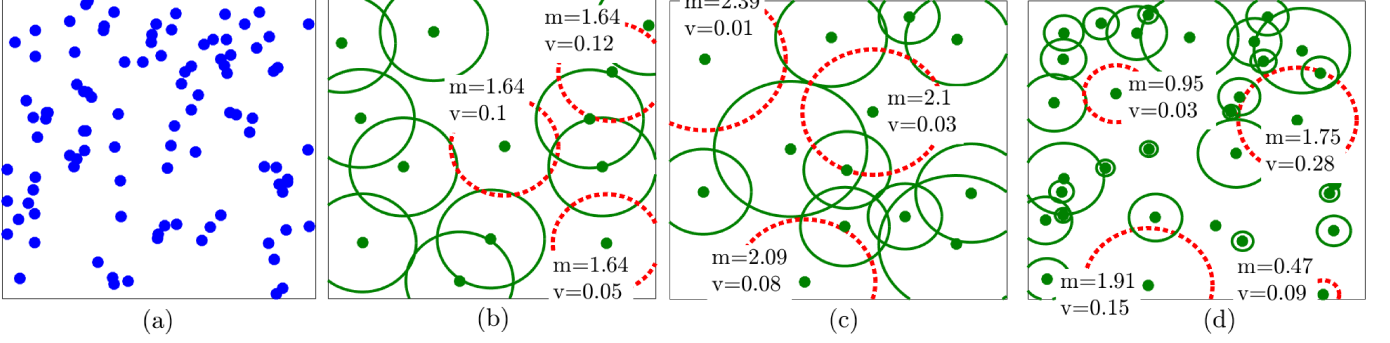


Fig. 1. SSSC p.p. realizations: (a) Begin with a realization of PPP Φ . Associate a gamma distributed mark $m \sim \gamma(\alpha, \beta)$ to each point where $\bar{m} = \alpha\beta$ is the average mark and $\alpha\beta^2$ is its variance. As β increases the mark variance increases. Associate a weight $v \sim U[0, 1]$ to each node independently. A node $x \in \Phi$ is selected based on (4), where $p_0 = 1$, $f_c(r, m, n)$ as in (5) with $c = 100$. The marks and weights for some retained points are shown in dotted circles. The retained p.p.s for (b) $\beta = 0$ (fixed mark radii), (c) $\beta = 0.1$ (mark radii have low variance), and (d) $\beta = 1$ (mark radii have high variance). As β increases, packing is denser, which is desired for spatially balanced caching.

We next detail the dependent thinning procedure, and investigate the relationship between Φ and $\Phi_{th,i}$, $i = \{1, \dots, M\}$.

A. Dependent Sampling of Nodes for Placement

In this section and onwards, for brevity of notation, we omit the index i , and consider the generic thinned process Φ_{th} , which is derived from Φ by applying the following probabilistic dependent thinning rule. Assume that mark m has a distribution μ , and $\tilde{\delta} = \{m\}$ is the set of marks for all points in $\tilde{\Phi}$, where $m \sim \mu$ and $\bar{m} = \mathbb{E}[m]$. Assume that weight ν_m does not depend on the mark m . The marked point $(x, m, v) \in \tilde{\Phi}$ is retained as a point of Φ_{th} with probability

$$p(x, m, v, \Phi) = p_0 \prod_{(y, n, w) \in \Phi, y \neq x} [1 - \mathbb{1}\{v \geq w\} f(\|x - y\|, m, n)] \quad (4)$$

independently from deleting or retaining other points of Φ . In other words, a node $x \in \Phi$ is retained to cache item i with probability p_0 , if it has the lowest weight among all the points within its exclusion range. In (4), $p_0 \in (0, 1]$, $f : [0, \infty[\mathbb{R}^2 \rightarrow [0, 1]$ is a deterministic function satisfying $f(\cdot, m, n) = f(\cdot, n, m)$ for all $m, n \in \mathbb{R}$. This means that if two points with marks m and n , and weights $v \geq w$ are a distance $r > 0$ apart, then the point with weight v is deleted by the other point with probability $f(r, m, n)$. Additionally, each surviving point is then again independently p_0 -thinned. The function $f(\|x - y\|, m, n)$ in (4) should be determined according to (3). Inspired from [17], assume that f satisfies

$$f_c(r, m, n) = \exp(-c[r - m - n]_+), \quad r \geq 0. \quad (5)$$

Denote by $\text{SSCC}[\lambda, \mu, (\nu_m)_{m \in \mathbb{R}}, p_0, f]$ the distribution of Φ_{th} . We next give its intensity, i.e., $\lambda_{th} = \lambda \mathbb{E}[p(x, m, v, \Phi)]$.

Theorem 1. [17, Theorem 12] *The intensity of the process $\Phi_{th} \sim \text{SSCC}[\lambda, \mu, (\nu_m)_{m \in \mathbb{R}}, p_0, f]$ is given by*

$$\lambda_{th} = \lambda p_0 \int_{\mathbb{R}} \int_{\mathbb{R}} \exp\left(-\lambda \int_{\mathbb{R}} F_{\nu_n}(w) \int_{\mathbb{R}^2} f(\|x\|, m, n) dx \mu(dn)\right) \nu_m(dw) \mu(dm),$$

where F_{ν_m} is the cumulative distribution function of ν_m .

Proof. The probability generating functional (PGFL) [18] of the PPP states for function $f(x)$ that $\mathbb{E}[\prod_{x \in \Phi} f(x)] = \exp(-\lambda \int_{\mathbb{R}^2} (1 - f(x)) dx)$. We obtain λ_{th} using the PGFL and $\mathbb{E}[\mathbb{1}\{\nu_n \leq w\}] = \int_{\mathbb{R}} F_{\nu_n}(w) \mu(dn)$, along with (4). \square

In Fig. 1, we plot different realizations of SSSC Φ_{th} formed by thinning Φ . As the mark variance increases, the packing is denser, which is desired for spatially balanced caching.

B. Spherical Contact Distribution Function

Our goal in this section is to relate the cache hit probability distribution to the (spherical) contact distribution function.

Definition 1. *The spherical contact distribution function (SCDF) of the p.p. Ξ is the conditional distribution function of the distance from a point chosen randomly outside Ξ (i.e. 0), to the nearest point of Ξ given $0 \notin \Xi$ [18]. It is given by*

$$H(r) = \mathbb{P}(R_{\text{Sph}} \leq r | R_{\text{Sph}} > 0), \quad r \geq 0, \quad (6)$$

where $R_{\text{Sph}} = \inf\{s : \Xi \cap sA \neq \emptyset\}$, where $A = B_0(1)$, and rA is the dilation of the set A by the factor r .

As an example, in Fig. 2 we show the SCDF for the Boolean model with random spherical grains in [19, Ch. 3.1].

Theorem 2. *The average cache hit probability of policy π is*

$$\mathbb{E}_{\pi}[F(Z)] = \mathbb{E}_{\mathcal{I}}[H_{\pi, \mathcal{I}}(R_c)], \quad (7)$$

where $H_{\pi, \mathcal{I}}(R_c)$ is the SCDF of the thinned p.p. $\Phi_{th, \mathcal{I}}$ for \mathcal{I} .

Proof. Let $B = B_0(R_c)$ and $\Phi_{th, i}(B) = \sum_{x \in \Phi_{th, i}} \mathbb{1}(x \in B)$ be the number of transmitters containing item i within a circular region of radius R_c around the origin. Then we have

$$F(Z) = \sum_i p_r(i) \mathbb{1}(\Phi_{th, i}(B) > 0).$$

The average cache hit probability is given by $\mathbb{E}[F(Z)] = \sum_i p_r(i) \mathbb{P}(\Phi_{th,i}(B) > 0)$, where defining $R_{Sph} = \inf\{s : \Phi_{th,i}(B_0(s)) \neq \emptyset\}$, given $0 \notin \Phi_{th,i}$ we have that

$$\mathbb{P}(\Phi_{th,i}(B) > 0) = \mathbb{P}(R_{Sph} \leq R_c | R_{Sph} > 0), \quad (8)$$

which is the SCDF of $\Phi_{th,i}$ evaluated at R_c . \square

The variance of $F(Z)$ across the nodes satisfies

$$\text{Var}_\pi[F(Z)] = \sum_i p_r^2(i) H_{\pi,i}(R_c)(1 - H_{\pi,i}(R_c)) \quad (9)$$

since the spatial thinning processes across different items are independent of each other. Under the IRM and a Zipf popularity model, $\text{Var}_\pi[F(Z)]$ decreases with increasing variance of marks when $\mathbb{E}[C(x)]$ is held constant. A spatially balanced sampling yields a low $\text{Var}_\pi[F(Z)]$ as expected.

C. Migration to the Child Process: Effective Thinning

Consider the pair $\Phi - \Phi_{th}$ of mother and child p.p.'s. The spherical contact distance denotes the distance between a typical point in Φ and its nearest neighbor from Φ_{th} .

Using (6), the SCDF for the p.p. Φ can be written as:

$$H_\pi(R) = 1 - \exp\left(-\int_0^R 2\pi r \lambda \eta_\pi(r, \delta) dr\right), \quad (10)$$

where $\eta_\pi(r, \delta)$ is the conditional thinning Palm-probability (CTPP), i.e. the probability of the point $x \in \Phi$ migrating to Φ_{th} under policy π , with a fixed (exclusion) radius δ . It equals

$$\eta_\pi(r, \delta) = \mathbb{P}(x \in \Phi_{th} | \Phi_{th} \cap B_{x_0}(r) = \emptyset, x_0 \in \Phi). \quad (11)$$

Remark 1. An effective thinning policy yields a larger CTPP $\eta_\pi(r, \delta)$. The more effective the thinning is, the larger (10) is. From Theorem 2, $\mathbb{E}_\pi[F(Z)]$ is improved if π is more effective.

We next compute the CTPP for the SSCC policy.

Proposition 2. The CTPP for PPP-SSCC is given as

$$\eta_{SSCC}(r, \vec{\delta}) = \int_{\mathbb{R}} \int_0^1 e^{-u\lambda \int_{\mathbb{R}^2} h(\|x\|, m, n) dx \mu(dn)} du \mu(dm),$$

where given radius marks m, n , $h(\|x\|, m, n)$ satisfies the relation $\int_{\mathbb{R}^2} h(\|x\|, m, n) dx = \pi(m+n)^2 - l_2(r, n)$, where $l_2(r, \delta)$ is the area of the intersection of $B_{x_0}(r)$ and $B_x(\delta)$.

Proof. The proof follows from generalizing [20, Eq. (15)]. \square

Corollary 1. The CTPP for PPP-MatII is given as

$$\eta_{MatII}(r, \delta) = \frac{1 - e^{-\lambda(\pi\delta^2 - l_2(r, \delta))}}{\lambda(\pi\delta^2 - l_2(r, \delta))}.$$

The next Theorem shows that having a distribution on the marks yields a more effective thinning than MatII does.

Theorem 3. The CTPPs satisfy $\eta_{SSCC}(r, \vec{\delta}) \geq \eta_{MatII}(r, \bar{m})$, where $\vec{\delta} = \{m\}$ is the set of marks in Φ , with $\bar{m} = \mathbb{E}_m[m]$.




Fig. 2. The radius of the smallest sphere centered at 0 and intersecting the Boolean Model Φ_{th} . The SCDF is the conditional distribution function of the radius of the sphere, given $0 \notin \Phi_{th}$ [19, Ch. 3.1].

Proof. From Prop. 2, we have that

$$\begin{aligned} \eta_{SSCC}(r, \vec{\delta}) &= \int_{\mathbb{R}} \int_0^1 e^{-u\lambda \int_{\mathbb{R}} (\pi(m+n)^2 - l_2(r, n)) \mu(dn)} du \mu(dm) \\ &= \mathbb{E}_m \left[\mathbb{E}_U \left[e^{-Uq(\lambda, r, m)} \right] \right] = \mathbb{E}_m \left[\frac{1 - e^{-q(\lambda, r, m)}}{q(\lambda, r, m)} \right], \end{aligned}$$

where $U \sim U[0, 1]$, and $q(\lambda, r, m) = \lambda\pi(m^2 + 2m\bar{m}_2) + \lambda\mathbb{E}_{m_2}[\pi m_2^2 - l_2(r, m_2)]$. Let $f = e^{-x}$, $x = U\lambda\pi(m^2 + 2m\bar{m}_2)$. We have, $\eta_{SSCC}(r, \vec{\delta}) = \mathbb{E}_m[g(m)]$, with $g = \frac{f-1}{\log(f)} = \frac{1-e^{-x}}{x}$. Then $g' = \frac{e^{-x}(x+1)-1}{x^2}$, $g'' = \frac{2-e^{-x}(x^2+2x+2)}{x^3} > 0$. Hence, $\eta_{SSCC}(r, \vec{\delta}) = \mathbb{E}_m[g(m)] \geq g(\bar{m}) = \eta_{MatII}(r, \bar{m})$. \square

Exploiting Theorem 3, $\eta_{SSCC}(r, \vec{\delta})$ can be improved using a mixture of marks. The variable exclusion range model can suit to the case of cellular networks where demand is not uniform across the network [14], which we left as future work.

D. Cache Over-Utilization

The cache placement requires $C(x) = \sum_i \mathbb{1}_{x \in \Phi_{th,i}} \leq N$, for all $x \in \Phi$, where N is finite. The storage constraint is satisfied on average, i.e. $N = \mathbb{E}[C(x)] = \sum_i p(x, m_i, v, \Phi)$, $x \in \Phi$. However, the set of child p.p.'s $\{\Phi_{th,i}\}_i$, $i = 1, \dots, M$ might overlap. We need to make sure that the cache capacities are not over-utilized. Hence, the intersection of the sampled processes, i.e. $\cap_i \Phi_{th,i}$, should not include any $x \in \Phi$ more than N times with high probability. We next provide an upper bound for the violation probability of the cache size for SSCC.

Proposition 3. Bernstein bound for cache size. The cache violation probability is upper bounded as

$$\mathbb{P}(C(x) > C) \leq \exp\left(-\frac{(C - N)^2}{\text{Var}[C(x)] + \frac{1}{3}(C - N)}\right), \quad (12)$$

where $\text{Var}[C(x)] = \sum_{i=1}^M \text{Var}[z_{xi}]$ since the placement is independent across items, where $\text{Var}[z_{xi}] = \mathbb{E}[z_{xi}^2] - \mathbb{E}[z_{xi}]^2 = p(x, m_i, v, \Phi)(1 - p(x, m_i, v, \Phi))$, for $i \in \{1, \dots, M\}$, $x \in \Phi$.

Proof. It follows from employing Bernstein inequality since $\mathbb{1}_{x \in \Phi_{th,i}}$ are independent 1-0 random variables across i . \square

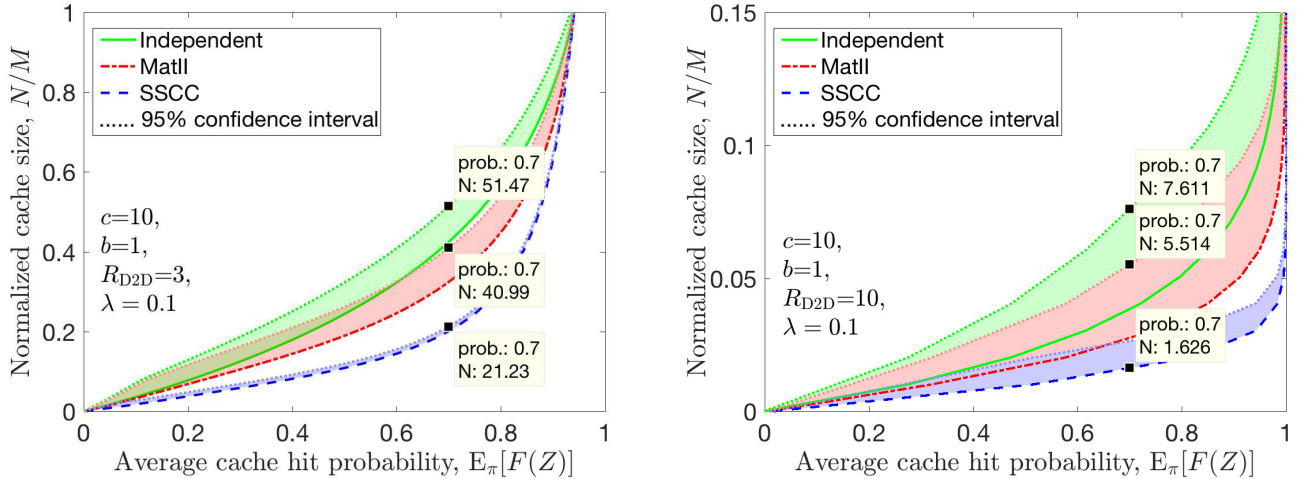


Fig. 3. The normalized cache size versus average cache hit rate for different placement policies. (Left) $R_c = 3$, (Right) $R_c = 10$.

As $\text{Var}[C(x)]$ drops, the bound in (12) becomes lower. Hence, the cache violation probability is negligible if the cache placement strategy has very low-variance. In Sect. IV, we demonstrate that SSCC has very small violation probability.

For the spatially independent placement policy in [11], where nodes are sampled i.i.d., authors have proposed a probabilistic placement technique to guarantee that the cache constraint is satisfied with equality. However, in SSCC, nodes are not sampled independently. Because the placement policy is NA across the nodes, it is nontrivial to design probabilistic placement techniques to satisfy the cache size constraint. In this section, we discuss how to bound the violation probability, and demonstrate in Sect. IV that for SSCC the cache violation probability can be made negligibly small.

IV. NUMERICAL SIMULATIONS

The nodes live in a square region of the Euclidean plane with area L^2 where $L = 100$. To avoid edge effects, we evaluate the performance only for the middle square region with area $L^2/9$. The network parameters are $\lambda = 0.1$ and $R_c \in \{3, 10\}$. The request process is isotropic and Zipf distributed with parameter $\gamma_r = 0.1$ over $M = 100$ items.

For MatII, there is a fixed exclusion range for a given item, and we have derived the optimal exclusion radii in [13]. Let r_i be the optimized exclusion range for item i for MatII. For SSCC, we assume that the marks $m^{(i)}$ for item i (exclusion radii) are distributed according to a gamma distribution $\mu^{(i)} = \Gamma(0.7r_i, 1)$ for each $x \in \Phi$, and all items i , where we choose its parameters such that the average value of the radius mark for item i equals $\bar{m}^{(i)} = 0.7r_i$. Hence, $\Phi_{th} \sim \text{SSCC}[0.1, \Gamma(0.7r_i, 1), U[0, 1], 1, f_{10}]$. We can observe that the SSCC model can be used to optimize the cache hit probability-cache violation probability tradeoff. As variance of exclusion range increases, the violation probability might also increase for a desired cache hit probability. Note that we do not optimize the distributions of the marks $\mu^{(i)}$ across all i over a

class of distributions. We leave the study of the fundamental performance limits of SSCC as future work.

We numerically investigate how much cache over-provisioning is required for different spatial cache placement policies: spatially independent [11], MatII [13], and SSCC cache placement. In Fig. 3, we investigate the required cache size N (normalized) of each policy given that the probability of cache violation is small such that $\mathbb{P}[|C(x) - N| \leq \epsilon] > 0.95$ in order to characterize the required cache size for a given average cache hit probability. We also illustrate the 95% confidence intervals represented by the shaded regions, and mark the cache sizes for different policies when the average cache hit probability is $E_\pi[F(Z)] = 0.7$. For example, when $R_c = 3$, for the 95% confidence interval, the excess cache ratio for independent placement in [11], and MatII placement in [13] with respect to the SSCC policy is 142%, and 93%, respectively. When we have $R_c = 10$, the respective excess ratios for the independent and MatII placement policies are 188%, and 109%, which are illustrated on the plots. SSCC yields a better concentration of the required cache size, which is desired. Hence, policies like SSCC can be exploited so that the cache does not overrun or underrun its capacity constraint.

SSCC gives insights into not only how to cache the content, but also how to effectively sample in spatial settings. SSCC is suited for enabling applications such as D2D and P2P as it promotes the item diversity and reciprocation. Extensions include the incorporation of the spatial variation of the demand. They also include employing the exclusion based models to optimize the performance of time-to-live (TTL) caches.

ACKNOWLEDGMENT

We thank Salman Salamatian for helpful discussions.

REFERENCES

- [1] N. G. K. Shanmugam, A. G. Dimakis, A. F. Molisch, and G. Caire, "FemtoCaching: Wireless content delivery through distributed caching helpers," *IEEE Trans. Inf. Theory*, vol. 59, no. 12, pp. 8402–13, Dec. 2013.

- [2] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Trans. Inf. Theory*, vol. 60, no. 5, pp. 2856–67, May 2014.
- [3] Q. Yu, M. A. Maddah-Ali, and A. S. Avestimehr, "The exact rate-memory tradeoff for caching with uncoded prefetching," *IEEE Trans. Inf. Theory*, vol. 64, no. 2, pp. 1281–1296, 2018.
- [4] M. Ji, G. Caire, and A. F. Molisch, "Fundamental limits of caching in wireless D2D networks," *IEEE Trans. Inf. Theory*, vol. 62, no. 2, pp. 849–869, Feb. 2016.
- [5] M. A. Maddah-Ali and U. Niesen, "Cache-aided interference channels," in *Proc., IEEE Int. Sym. Inf. Theory*. IEEE, 2015, pp. 809–813.
- [6] S.-H. Park, O. Simeone, and S. Shamai, "Joint optimization of cloud and edge processing for fog radio access networks," in *Proc., IEEE Int. Sym. Inf. Theory*, 2016, pp. 315–319.
- [7] H. Che, Y. Tung, and Z. Wang, "Hierarchical web caching systems: Modeling, design and experimental results," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 7, pp. 1305–14, Sep. 2002.
- [8] A. Giovanidis and A. Avranas, "Spatial multi-LRU caching for wireless networks with coverage overlaps," in *Proc., ACM Sigmetrics/IFIP Performance*, Antibes, France, Jun. 2016, pp. 403–405.
- [9] D. S. Berger, S. Henningsen, F. Ciucu, and J. B. Schmitt, "Maximizing cache hit ratios by variance reduction," in *ACM Sigmetrics Performance Evaluation Review*, vol. 43, no. 2, Sep. 2015, pp. 57–59.
- [10] A. Legout, G. Urvoy-Keller, and P. Michiardi, "Rarest first and choke algorithms are enough," in *Proc., ACM Sigcomm*, Oct. 2006, pp. 203–216.
- [11] B. Błaszczyszyn and A. Giovanidis, "Optimal geographic caching in cellular networks," in *Proc., IEEE ICC*, UK, Jun. 2015, pp. 3358–3363.
- [12] S. Ioannidis and E. M. Yeh, "Adaptive caching networks with optimality guarantees," in *Proc., ACM SIGMETRICS*, Jun. 2016, pp. 113 – 124.
- [13] D. Malak, M. Al-Shalash, and J. G. Andrews, "Spatially correlated content caching for device-to-device communications," *IEEE Trans. Wireless Commun.*, vol. 17, no. 1, pp. 56–70, Jan. 2018.
- [14] S. Zhou, D. Lee, B. Leng, X. Zhou, H. Zhang, and Z. Niu, "On the spatial distribution of base stations and its relation to the traffic density in cellular networks," *IEEE Access*, vol. 3, pp. 998–1010, 2015.
- [15] S. Traverso, M. Ahmed, M. Garetto, P. Giaccone, E. Leonardi, and S. Niccolini, "Temporal locality in today's content caching: why it matters and how to model it," *ACM Sigcomm Computer Commun. Review*, vol. 43, no. 5, pp. 5–12, 2013.
- [16] D. P. Dubhashi and D. Ranjan, "Balls and bins: A study in negative dependence," *BRICS Report Series*, vol. 3, 1996.
- [17] J. Teichmann, F. Ballani, and K. G. van den Boogaart, "Generalizations of Matérn's hard-core point processes," *Spatial Statistics*, vol. 3, pp. 33–53, 2013.
- [18] D. Stoyan, W. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*, 2nd ed. John Wiley and Sons, 1996.
- [19] F. Baccelli and B. Błaszczyszyn, *Stochastic Geometry and Wireless Networks*. NOW: Found. Trends. Network., 2010.
- [20] A. Al-Hourani, R. J. Evans, and S. Kandeepan, "Nearest neighbor distance distribution in hard-core point processes," *IEEE Communications Letters*, vol. 20, no. 9, pp. 1872–1875, 2016.